

APPLICATION FOR UNITED STATES PATENT

INVENTORS

OPHIR FRIEDER
ABDUR R. CHOWDHURY

TITLE:

SYSTEM FOR SIMILAR DOCUMENT DETECTION

ATTORNEYS' ADDRESS:

VENABLE
1201 New York Avenue, N.W., Suite 1000
Washington, D.C. 20005-3917
Telephone: (202) 962-4800
Telefax: (202) 962-8300

ADDRESS FOR U.S.P.T.O. CORRESPONDENCE:

VENABLE
Post Office Box 34385
Washington, D.C. 20043-9998

ATTORNEY DOCKET NO.:

7519-164345

SYSTEM FOR SIMILAR DOCUMENT DETECTION

BACKGROUND OF THE INVENTION

Field of the Invention

5 The invention relates to a system for similar document detection and, more particularly, to detecting whether a document is similar to another document in a document collection.

References

10 For the convenience of the reader, the publications referred to in the specification are listed below. In the specification, the numerals within square brackets refer to the respective 15 publications. The listed publications are incorporated herein by reference.

[1] L. Giles and S. Lawrence, "Accessibility and Distribution of Information on the Web," Nature, vol. 400, 1999.

[2] http://nccam.nih.gov, The National Institutes of Health (NIH), National Center for Complementary and Alternative Medicine (NCCAM), April 12, 2000.

[3] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig, "Syntactic Clustering of the Web," Sixth International World Wide Web Conference, April, 1997.

[4] S. Brin, J. Davis, and H. Garcia-Molina, "Copy Detection Mechanisms for Digital Documents," Proceeding of SIGMOD '95, 1995.

[5] N. Shivakumar and H. Garcia-Molina, "Finding Near-Replicas of Documents on the Web," Proceedings of Workshop on Web Databases (WebDB'98), March 1998.

[6] N. Shrivakumar and H. Garcia-Molina, "SCAM: A Copy Detection Mechanism for Digital Documents," Proceedings of the Second International Conference on Theory and Practice of Digital Libraries, 1995.

[7] N. Shrivakumar and H. Garcia-Molina, "Building a Scalable and Accurate Copy Detection Mechanism," Proceedings of Third International Conference on Theory and Practice of Digital Libraries, 1996.

[8] N. Heintze, "Scalable Document Fingerprinting," Proceedings of the Second USENIX Workshop on Electronic Commerce, 1996.

[9] C. Buckley, C. Cardie, S. Mardis, M. Mitra, D. Pierce, K. Wagstaff, and J. Walz, "The Smart/Empire TIPSTER IR System," TIPSTER Phase III Proceedings, Morgan Kaufmann, 2000.

Selv A
[10] V. Chalana, A. Bruce, and T. Nguyen, "Mathsoft Data Analysis Products Division: DocBrowser," www.statsci.com/docbrowse/paper/spie98/node1.htm, July 31, 1999.

Selv A
[11] G. Salton, C.S. Yang, and A. Wong, "A Vector-Space Model for Information Retrieval," Comm. of the ACM, vol. 18, 1975.

Selv A
[12] M. Porter, "An Algorithm for Suffix Stripping," Program, vol. 14, no. 3, pp. 130-137, 1980.

10 [13] B. Kjell, W.A. Woods, and O. Frieder, "Discrimination of Authorship Using Visualization," Information Processing and Management, Pergamon Press, vol. 30, no. 1, pp. 141-150, January 1994.

Selv A
[14] R. Scotti and C. Lilly, George Washington University Declassification Productivity Research Center, <http://dprc.seas.gwu.edu>, July 31, 1999.

15 *Selv A*
[15] D. Grossman, D. Holmes, and O. Frieder, "A DBMS Approach to IR in TREC-4," Overview of the Fourth Text Retrieval Conference (TREC-4), November 1995.

20 [16] A. Smeaton, F. Kelley, and G. Quinn, "Ad Hoc Retrieval Using Thresholds, WSTs for French Monolingual Retrieval, Document-at-a-Glance for High Precision and Triphone Windows for Spoken Documents," Proceedings of the Sixth Text Retrieval Conference (TREC-6), p.461, 1997.

25 [17] U.S. Department of Commerce, National Institute of Standards and Technology, "Secure Hash Standard," Federal Information Processing Standards Publication FIPS PUB 180-1, April 17, 1995.

25 [18] U.S. Department of Commerce, National Institute of Standards and Technology, "Secure Hash Standard," Federal Information Processing Standards Publication FIPS PUB 180, May 11, 1993.

[19] Ronald L. Rivest, "The MD4 Message Digest Algorithm," Proceedings of Advances in Cryptology - CRYPTO '90, Springer-Verlag, pp. 303-311, 1991.

Background of the Invention

The detection of whether a document is similar to another document in a document collection is becoming an important problem due to the tremendous growth of the Internet and data portals (see, e.g., [1]). Document collections are increasing in both the number of 5 documents collections and in the number of documents in document collections due to the ease of transmitting and receiving documents via the Internet and data portals. As the size of a document collection increases, the probability of similar documents being re-submitted or re-indexed in the document collection increases as well. Maintaining similar documents in a document collection not only drains valuable resources for the computation and storage of 10 indices for the document collection but also affects the collection statistics and, hence, potentially the accuracy of searching the document collection.

Storing similar documents in a document collection affects both the accuracy and efficiency of an information search and retrieval engine used with the document collection. 15 Retrieving similar documents in response to a user's query potentially lowers the number of valid responses provided to the user, which thereby lowers the accuracy of the user's response set. Further, processing similar documents necessitates additional computation without introducing any additional benefit to the user, which lowers the processing efficiency of the user's query.

Additionally, similar documents skew collection statistics of the document collection. 20 Collection statistics are typically used as part of a similarity computation of a query for the document collection. With similar documents in the document collection, the collection statistics of the document collection are biased and may affect the overall precision of the document collection and its information search and retrieval engine.

The need to detect similar documents arises in various types of document collections. As 25 an example, for a document collection of documents received via the Internet, similar documents are undesirable additions to the document collection, and similar document detection would be useful prior to adding another document to the document collection. As another example, for a document collection of classified documents, similar documents need to be identified for either declassification or classification, and similar document detection would be useful for document 30 declassification or document classification. As another example, for a document collection of electronic mail (e-mail) documents, similar e-mail documents need to be identified for

processing the document collection, and similar document detection would be useful for e-mail traffic processing, which may be continuous. As a potential difficulty with processing a document collection of e-mail documents, many e-mail documents have a short length and, hence, may prove difficult to detect as similar documents.

5 As another example of the need to detect similar documents, consider searching web documents (e.g., documents available over the Internet via the world wide web format), which typically have a short length (e.g., typically around 4 kilobytes (KB) (see, e.g., [1])). With web documents, one might believe that matching the uniform resource locator (URL) would identify similar documents. However, because many web sites use dynamic presentation, where the
10 content changes depending on the region or other variables, relying on the URL is of little value. Further, data providers often create names for one web site to attract users with different interests or perspectives. For example, the web sites www.fox4.com, onsale.channel9.com, and www.realtv.com all point to an advertisement for realTV.

15 As another example of the need for detection of similar documents, similar documents can populate a document collection when multiple document sources are used. For instance, the National Center for Complementary and Alternative Medicine (NCCAM) (see, e.g., [2]),
20 supports an information search and retrieval engine for a document collection of medical data having inputs from multiple sources of medical data. Given the nature of the medical data, similar documents in the document collection can be common. Because unique document identifiers are most likely not possible when the document identifiers originate from different sources, the detection of similar documents is essential to produce non-redundant results for the
25 information search and retrieval engine.

Conventional techniques for detecting similar documents can be divided into three categories: shingling techniques; similarity measure techniques; and image processing
25 techniques. As the first category, shingling techniques were developed by various researchers, for example: DSC [3], COPS [4]; SCAM [5], [6], and [7], which is a successor to COPS; and KOALA [8]. The shingling technique uses a set of contiguous terms, or shingles, for a document and compares the number of matching shingles. The shingles can be considered to be subdocuments for the document. With the comparison of subdocuments between two
30 documents, a percentage of overlap is calculated between the two documents. For the shingling technique, a hash value is determined for each subdocument, and the hash values for each

subdocument are filtered to reduce the number of comparisons performed, which improves the runtime performance of the shingling technique. With the shingling technique, a document is not compared to other documents, and instead, subdocuments are compared. By using subdocuments, instead of documents, each comparison may produce many potential similar documents. With the shingling technique, by returning many potential matches of similar documents, a large amount of user involvement is required to sort potential similar documents, which dilutes the usefulness of the shingling technique.

To overcome the basic efficiency issues with the shingling technique, several optimization techniques for the shingling technique were proposed to reduce the number of comparisons made. For example, removing frequently occurring shingles (see, e.g., [8]) and retaining only every twenty-fifth single (see, e.g., [3]) were proposed. With these optimization techniques, the computation time of the shingling technique is reduced. However, because no semantic premise is used to reduce the volume of data, a degree of randomness is introduced to the comparison process, which results in relatively non-similar documents being identified as potential similar documents.

In terms of computational time complexity, the shingling technique has order $O(kd \log(kd))$, where k is the number of shingles per document, and d is the number of documents in the document collection. Even with the performance-improving technique of removing shingles occurring in over 1,000 documents and keeping only every twenty-fifth shingle, the implementation of the DSC took 10 days to process 30 million documents [3].

As an alternative to the DSC shingling technique, the DSC-SS shingling technique was proposed [3]. The DSC-SS shingling technique uses super shingles, in which several shingles are combined in a super shingle, which results in a document having a few super shingles, instead of many shingles. With the DSC shingling technique, the similarity between documents was measured as a ratio of matching shingles in two documents, and with the DSC-SS shingling technique, the similarity between two documents is measured using one super shingle for each document. Using a single super shingle is more efficient than using multiple shingles because a full counting of all overlaps between shingles is no longer required.

In terms of computational time complexity, the run time for DSC-SS shingling technique is of the order $O(kd \log(kd))$, although k is significantly smaller for the DSC-SS shingling technique compared to the k for the DSC shingling technique. Further, the amount of

computations required to count overlap is eliminated with the DSC-SS shingling technique, which reduces the overall runtime. Nonetheless, the DSC-SS shingling technique reportedly does not work well for documents having a short length. Moreover, the shingling technique and its optimization attempts are very sensitive to adjustments in the size of shingles and the
5 frequency of retained shingles.

As the second category, similarity measure techniques were proposed in, for example, [9] and [10]. The similarity measure techniques are similar to prior work done in document clustering (see, e.g., [11]). A similarity measure technique uses similarity computations to group potentially similar documents and compares each document pair-wise. Because of the pair-wise
10 comparison, a similarity measure technique is computationally prohibitive because the computational time complexity is of the order $O(d^2)$, where d is the number of documents.

To make the similarity measure technique computationally feasible, document terms are identified for each document, and a document collection is searched using the document terms. With the enhanced similarity measure technique, document terms are initially identified for a document to be compared to the document collection. Each term for the document is used to search the document collection, and a final weight is produced for each document in the document collection having a matching term. The document in the document collection having the largest weight is determined to be the most similar document. By using the document as a query with the similarity measure technique, clustering of the documents results. Even the enhanced similarity measure technique becomes computationally unfeasible for a large or dynamic document collection because each document must be queried against the entire collection.
20

As the third category, image processing techniques were proposed in, for example, [13] and [14]. An image processing technique processes documents as images to determine similar
25 documents. The image processing technique maps the similar document detection problem into an image-processing domain, rather than into the text-processing domain as with the shingling technique, the similarity measure technique, and the parsing filtering technique.

There exists a need for a technique to detect whether a document is similar to another document in a document collection, where the technique is scalable to and computationally
30 feasible for any size of document and any size of document collection.

SUMMARY OF THE INVENTION

An object of the invention is to provide a technique to detect whether a document is similar to another document in a document collection, where the technique is scalable to and computationally feasible for any size of document and any size of document collection.

5 The invention includes a method, an apparatus, and an article of manufacture for detecting similar documents.

The method of the invention includes a method for detecting similar documents comprising the steps of: obtaining a document; filtering the document to obtain a filtered document; determining a document identifier for the filtered document and a hash value for the 10 filtered document; generating a tuple for the filtered document, the tuple comprising the document identifier for the filtered document and the hash value for the filtered document; comparing the tuple for the filtered document with a document storage structure comprising a plurality of tuples, each tuple in the plurality of tuples representing one of a plurality of documents, each tuple in the plurality of tuples comprising a document identifier and a hash value; and determining if the tuple for the filtered document is clustered with another tuple in the 15 document storage structure, thereby detecting if the document is similar to another document represented by the another tuple in the document storage structure.

The method of the invention includes a method for detecting similar documents comprising the steps of: obtaining a document; parsing the document to remove formatting and to obtain a token stream, the token stream comprising a plurality of tokens; retaining only 20 retained tokens in the token stream by using at least one token threshold; arranging the retained tokens to obtain an arranged token stream; processing in turn each retained token in the arranged token stream using a hash algorithm to obtain a hash value for the document; generating a document identifier for the document; forming a tuple for the document, the tuple comprising the document identifier for the document and the hash value for the document; inserting the tuple for the document into a document storage tree, the document storage tree comprising a plurality of tuples, each tuple located at a bucket of the document storage tree, each tuple in the plurality of tuples representing one of a plurality of documents, each tuple in the plurality of tuples comprising a document identifier and a hash value; and determining if the tuple for the document 25 is co-located with another tuple at a same bucket in the document storage tree, thereby detecting 30

if the document is similar to another document represented by the another tuple in the document storage tree.

The method of the invention includes a method for detecting similar documents comprising the steps of: determining a hash value for a document; accessing a document storage structure comprising a plurality of hash values, each hash value in the plurality of hash values representing one of a plurality of documents; and determining if the hash value for the document is equivalent to another hash value in the document storage structure, thereby detecting if the document is similar to another document represented by the another hash value in the document storage structure.

10 The method of the invention includes a method for detecting similar documents comprising the step of: comparing a document to a plurality of documents in a document collection using a hash algorithm and collection statistics to detect if the document is similar to any of the documents in the document collection.

15 The system of the invention includes a computer system including a computer-readable medium having software to operate a computer in accordance with the invention.

The apparatus of the invention includes a computer including a computer-readable medium having software to operate the computer in accordance with the invention.

The article of manufacture of the invention includes a computer-readable medium having software to operate a computer in accordance with the invention.

20 Moreover, the above objects and advantages of the invention are illustrative, and not exhaustive, of those which can be achieved by the invention. Thus, these and other objects and advantages of the invention will be apparent from the description herein, both as embodied herein and as modified in view of any variations which will be apparent to those skilled in the art.

25

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention are explained in greater detail by way of the drawings, where the same reference numerals refer to the same features.

Figure 1 illustrates a flow diagram for the invention.

30 Figure 2 illustrates a flow diagram for filtering a document.

Figure 3 illustrates a flow diagram for retaining only retained tokens in a token stream.

Figure 4 illustrates a flow diagram for generating a document storage structure for a document collection.

Figure 5 illustrates an embodiment of a document storage structure of the invention.

Figure 6 illustrates a plan view for a computer system for the invention.

5

DETAILED DESCRIPTION OF THE INVENTION

In describing the invention, the following definitions are applicable throughout.

A "document" refers to a computer-readable file comprising text. As an option, the document comprises text and at least one of the following: an image, a graphical object, audio, video, and the like. The invention is not restricted as to the size of a document.

A document is referred to as being "similar" to another document if the two documents contain roughly the same semantic content, whether or not the two documents are a precise syntactic match. For example, two documents are similar if the two documents are highly the same, near duplicates, or duplicates. As another example, two documents are similar if the two documents have a sufficient measure of resemblance (see, e.g., [3] and [8]).

A "token" refers to any representation of context in a document. Examples of a token include: a word; a phrase (e.g., "New York," or "to be or not to be"); a portion of a word (e.g., ing, or qu), which is known as an n-gram, or a bi-gram, a tri-gram, and so on; a symbol; a number; any plural thereof; and any combination thereof.

A "tuple" refers to a data structure having two or more portions.

A "computer" refers to any apparatus that is capable of accepting a structured input, processing the structured input according to prescribed rules, and producing results of the processing as output. Examples of a computer include: a computer; a general purpose computer; a supercomputer; a mainframe; a super mini-computer; a mini-computer; a workstation; a micro-computer; a server; an interactive television; and a hybrid combination of a computer and an interactive television. A computer can have a single processor or multiple processors, which can operate in parallel and/or not in parallel. A computer also refers to two or more computers connected together via a network for transmitting or receiving information between the computers. An example of such a computer includes a distributed computer system for processing information via computers linked by a network.

10

15

20

25

30

A "computer-readable medium" refers to any storage device used for storing data accessible by a computer. Examples of a computer-readable medium include: a magnetic hard disk; a floppy disk; an optical disk, such as a CD-ROM and a DVD; a magnetic tape; a memory chip; and a carrier wave used to carry computer-readable electronic data, such as those used in transmitting and receiving e-mail or in accessing a network.

"Software" refers to prescribed rules to operate a computer. Examples of software include: software; code segments; instructions; computer programs; and programmed logic.

A "computer system" refers to a system having a computer, where the computer includes a computer-readable medium embodying software to operate the computer.

A "network" refers to a number of computers and associated devices that are connected by communication facilities. A network involves permanent connections such as cables or temporary connections such as those made through telephone or other communication links. Examples of a network include: an internet, such as the Internet; an intranet; a local area network (LAN); a wide area network (WAN); and a combination of networks, such as an internet and an intranet.

Figure 1 illustrates a flow diagram for the invention. The flow diagram in Figure 1 is used to detect whether a document is similar to another document in a document collection. In one embodiment, the goal in the detection is to build a document collection without any similar documents. In another embodiment, the goal in the detection is to remove similar documents from the document collection. In another embodiment, the goal in the detection is to identify and/or group similar documents in the document collection. In another embodiment, the goal in the detection is to determine where similar documents in the document collection are stored.

In block 1, a document is obtained. In one embodiment, the document is a document in the document collection, and in another embodiment, the document is not a document in the document collection. The document is obtained through any manner used by a computer system to obtain a document. For example, the document is obtained from a computer-readable medium of the computer system. As another example, the document is obtained via a network, such as the Internet. As a further example, the document is obtained by scanning a paper document and using optical character recognition software and the scanned paper document to obtain the document.

In block 2, the document is filtered to obtain a filtered document. Preferably, the document is filtered using collection statistics. In one embodiment, the collection statistics are based on the document collection. In another embodiment, the collection statistics are based on knowledge of the language of the document. In yet another embodiment, the collection statistics 5 are based on one or more other document collections. Further, the collection statistics can be any combination of the above embodiments. Filtering the document for one embodiment of the invention is further discussed with respect to Figure 2.

Figure 2 illustrates a flow diagram for filtering the document to obtain the filtered document. In block 15, the document is parsed to obtain a token stream. The token stream of the document comprises at least one token. However, for a trivial case, a document having no tokens may result from the parsing.

In block 16, formatting is removed from the token stream. Examples of formatting in a document include font size, bold, and left justify. Formatting in a document affects whether a document is similar to another document. By removing formatting from the document, the number of documents that are considered similar is increased. In one embodiment, depending on the type of document, formatting in the document may be represented by format tags. If format tags are used, the format tags are removed from the token stream in block 16.

In block 17, the token stream is examined such that only retained tokens are retained in the token stream. A determination is made whether each token in the token stream should be retained (i.e., a retained token) or should be removed from the token stream. Those tokens in the token stream that are not retained as retained tokens are removed from the token stream. Generally, at least one token is determined not to be a retained token and is removed from the token stream. For the trivial case, all tokens are determined to be retained tokens, and no tokens are removed from the token stream. By retaining some tokens and removing other tokens from the token stream, the effects of non-content changes in the document are minimized.

In one embodiment, a token is removed from the token stream if the token appears more than once in the token stream. The token is removed from the token stream if the token is a duplicate of another token in the token stream. As an option, instead of removing a duplicate token from the token stream, a replication count is maintained for each token appearing more than once in the token stream.

SEARCHED
SERIALIZED
INDEXED
FILED

In another embodiment, a token is removed from the token stream if the token represents a stop word. A “stop word” is any non-discriminatory token within a given context. In the trivial case, the given context is a complete language. Examples of typical stop words in the English language include: is; a; an; and the. In a non-trivial case, the given context is a 5 document collection, and zero or more document collection-specific stop words are identified from the document collection using collection statistics.

In an additional embodiment, a token is removed from the token stream if the token represents a very frequent token and/or a very infrequent token. For example, a token is removed from the token stream if the token is a very frequent token and is retained in the token 10 stream as a retained token if the token is not a very frequent token. As another example, a token is removed from the token stream if the token is a very infrequent token and is retained in the token stream as a retained token if the token is not a very infrequent token. As yet another example, a token is retained in the token stream as a retained token if the token is neither a very frequent token nor a very infrequent token and is otherwise removed from the token stream. Because frequently occurring terms in a document often do not add to the semantic content of the 15 document (see, e.g., [15] and [16]), the invention removes very frequent tokens from the token stream. In addition to or in lieu of removing very frequent tokens from the token stream, the invention removes very infrequent tokens from the token stream. The collection statistics are used to determine whether a token is a very frequent token and/or a very infrequent token. The determination of the collection statistics, very frequent tokens, and very infrequent tokens are 20 further discussed below with respect to Figures 3 and 4.

In a further embodiment, a token is removed from the token stream based on a determination as to which part of speech the token represents. Examples of parts of speech include: a noun, a verb, an adjective, an adverb, a preposition; and a type of noun (e.g., a person, 25 a place, a thing). For example, a token is removed from the token stream if the token represents and/or does not represent a desired part of speech. As another example, a token is removed from the token stream if the token does not represent a noun, and a token is retained in the token stream as a retained token if the token represents a noun. As yet another example, a token is removed from the token stream if the token represents a preposition, and a token is retained in 30 the token stream as a retained token if the token does not represent a preposition.

In a still further embodiment, a token is removed from the token stream based on a ranking of a part of speech represented by the token. With the collection statistics, each token in the document collection is ranked according to its number of appearances in the document collection with respect to the part of speech represented by the token. For example, all tokens 5 representing nouns in the document collection are ranked as to their number of appearances in the document collection as compared to all other nouns in the document collection. With the rankings, a determination is made whether to remove or retain a token in the token stream. For example, a token is removed from the token stream if the token represents a noun that occurs very frequently in the document collection and is retained in the token stream as a retained token 10 if the token represents a noun that does not occur very frequently in the document collection. As another example, a token is removed from the token stream if the token represents a noun that occurs very infrequently in the document collection and is retained in the token stream as a retained token if the token represents a noun that does not occur very infrequently in the document collection. As yet another example, a token is retained in the token stream as a retained token if the token represents a noun that occurs neither very frequently nor very infrequently in the document collection and is otherwise removed from the token stream. 15

In addition to the above embodiments for block 17, other techniques are available for identifying which tokens to retain or remove from the token stream. Further, tokens can be retained or removed according to any combination of the above embodiments for block 17. The 20 above embodiment using very frequent tokens and/or very infrequent tokens to retain and remove tokens from the token stream is further discussed with respect to Figure 3.

Figure 3 illustrates a flow diagram for retaining only retained tokens in a token stream using identified very frequent tokens and/or very infrequent tokens. In this illustration, a token is retained as a retained token in a token stream if the token is neither a very frequent token nor a 25 very infrequent token. Two thresholds, namely threshold T1 and threshold T2, are determined based on the collection statistics for assessing whether a token is a very infrequent token or a very frequent token, respectively.

In block 21, a token in the token stream is obtained.

In block 22, a score is determined for the token. The score for the token is preferably 30 based on the collection statistics (e.g., the number of times a token occurs in the document collection and the number of documents in the document collection). For instance, an exemplary

equation for determining the score of a token is the equation for inverted document frequency (IDF):

$$\text{score} = \log(D/F(i)) \quad (1)$$

where D is the number of documents in the document collection and F(i) is the number of

5 documents in the document collection in which the token i appears.

In block 23, the score is compared to the threshold T1, which is the threshold used to identify if the token is a very infrequent token. If the score for the token is less than the threshold T1, flow proceeds to block 25. If the score for the token is greater than or equal to the threshold T1, flow proceeds to block 24.

10 In block 24, the score for the token is greater than or equal to the threshold T1, which indicates that the token is a very infrequent token. Hence, in block 24, the token is removed from the token stream. Flow proceeds to block 21 if there is another token in the token stream to process.

15 In block 25, the score for the token is less than the threshold T1, which indicates that the token is not a very infrequent token. In block 25, the score is compared to the threshold T2, which is the threshold used to identify a very frequent token. If the score for the token is less than the threshold T2, flow proceeds to block 24. If the score for the token is greater than or equal to the threshold T2, flow proceeds to block 26.

20 If the flow proceeds from block 25 to block 24, the token is removed from the token stream as being a very frequent token.

25 In block 26, the token is a token which is neither a very infrequent token nor a very frequent token, and the token is retained as a retained token in the token stream. From block 26, flow proceeds to block 21 if there is another token in the token stream to process.

As an option for Figure 3, the order of blocks 25 and 23 are switched. As another option, instead of a high score indicating a very frequent token, a high score indicates a very infrequent token. As another option, a single threshold is used by removing either block 23 or block 25 from Figure 3. As another option, other techniques for retaining and removing tokens from the token stream are included in Figure 3.

Referring back to Figure 2, in block 18, the retained tokens in the token stream are arranged to obtain an arranged token stream. Preferably, the retained tokens are arranged in Unicode ascending order. Unicode is a character and coding standard that represents almost all

of the written languages of the world. As an option, the retained tokens are arranged in Unicode descending order. As an option, the tokens are arranged according to any technique or standard. Examples of other techniques and standards for arranging the tokens include: alphabetizing; arranging by symbols; arranging by numbers; arranging by non-English letters; or arranging by another character and coding standard, such as the American Standard Code for Information Interchange (ASCII), extended ASCII, and Extended Binary Coded Decimal Interchange Code (EBCDIC).
5

Using the flow diagram of Figure 2 to implement block 2, the filtered document obtained from block 2 comprises an arranged token stream. Other techniques for obtaining a filtered 10 document can be used. For example, in one embodiment, the flow diagram of Figure 2 is implemented without block 18. As another example, the order of blocks 15 and 16 in Figure 2 are switched.

Referring back to Figure 1, in block 3, a document identifier is determined for the filtered 15 document. The document identifier uniquely identifies the document from the other documents in the document collection. For example, the document identifier can be a number which is serially determined based on a previous number assigned to a previous document.

In block 4, a hash value is determined for the filtered document. Using the arranged token stream for the filtered document from Figure 2, the hash value for the filtered document is determined by processing individually each retained token in the arranged token stream. 20 Examples of hash algorithms include the following: a hash algorithm; a standard hash algorithm; a secure hash algorithm; the hash algorithm SHA-1, which is described in [17]; the hash algorithm SHA, which is described [18]; the message digest algorithm MD 4, which is described in [19]; and any algorithm used to map an identifier to a numerical value associated with a location of a data structure.

25 Preferably, the hash value for the filtered document is determined using a hash algorithm having an approximately even distribution of hash values. More preferably, the hash value for the filtered document is determined using a secure hash algorithm. With a secure hash algorithm, the probability of two token streams creating the same hash value is reduced. Even more preferably, the hash value for the filtered document is determined using the hash algorithm 30 SHA-1. With the hash algorithm SHA-1, each retained token in the arranged token stream is processed individually to obtain a hash value for the filtered document. With the SHA-1 hash

algorithm, which uses 160 bits, the probability of duplicate values for different token streams is $P(2^{-160})$.

In block 5, a tuple is generated for the filtered document. Preferably, the tuple comprises the document identifier determined in block 3 and the hash value determined in block 4. As an option, additional information can be included. An example of additional information for the tuple includes the source of the document (e.g., the file pathname of the document).

In block 6, the tuple for the filtered document is compared with a document storage structure for the document collection. The document storage structure for the document collection comprises any data structure that is efficient for storing and accessing representations of the documents in the document collection. Examples of data structures useful for the document storage structure include the following: a hash table; a tree, such as a binary tree, a balanced binary tree, and a red-black binary tree; a queue; a ring; an array; and any combination of data structures. Preferably, the document storage structure is a combination of a hash table and a tree. More preferably, the document storage structure is a combination of a hash table and a binary tree. Even more preferably, the document storage structure is a combination of a hash table and a balanced binary tree.

Preferably, each document in the document collection is processed to obtain a tuple similar to the tuple generated in block 5, and each tuple representing a document in the document collection is stored in the document storage structure. The generation of the document storage structure for the document collection is further discussed below with respect to Figure 4.

Figure 5 illustrates a preferred embodiment of a document storage structure of the invention. The document storage structure 40 comprises a combination of a hash table 41 and at least one tree. The hash table 41 includes bins 42-1, 42-2, 42-3, etc. As an example, if the hash algorithm of block 4 is the hash algorithm SHA-1, the hash table has up to 2^{160} bins. Each bin 42-1, 42-2, 42-3, etc. has a corresponding tree 43-1, 43-2, 43-3, etc., respectively. Each tree 43-1, 43-2, 43-3, etc. is, for example, a balanced binary tree. Each bin has zero or more tuples from the document collection associated with the bin, and each tuple associated with the bin is arranged according to the tree corresponding to the bin. For example, bin 42-1 has a tuple 44 for the document collection attached to the tree 43-1. If a bin has no tuples, the tree is empty or comprises nulls, or the tree is not developed. If a bin has one or more tuples, the tree includes

the tuples. Each tuple for the document collection is hashed into the document storage structure 40.

To compare the tuple with the document storage structure, the tuple is preferably inserted into the document storage structure. Referring back to Figure 5, the tuple 45 for the document is 5 compared to the document storage structure 40. Preferably, the hash value for the tuple 45 is hashed into the hash table of the document storage structure 40. If bin 42-3 of the hash table 41 matches the hash value for the tuple 45, the tuple 45 is inserted into the document storage structure as tuple 45' in the tree 43-3 of bin 42-3. If a bin does not yet exist in the hash table for the tuple 45, a bin in the hash table and a corresponding tree are created for the tuple 45.

10 As an example, if the document storage structure is a hash table, the tuple for the filtered document is inserted into the hash table by placing the tuple into a bin (or, e.g., entry) of the hash table or by creating a bin for the hash table.

As another example, if the document storage structure is a tree, the tuple for the filtered document is inserted into the tree by inserting the tuple into an existing branch (or, e.g., bucket) of the tree or creating a new branch for the tree.

As an option, the hash value for the tuple, instead of the tuple, is inserted into the document storage structure.

As an option, the document storage structure is scanned to determine if a tuple exists which has a hash value that is the same as hash value for the tuple of the document.

20 As an option, the document storage structure is scanned to determined if a hash value exists which is the same as the hash value for the tuple of the document.

In block 7, it is determined whether the tuple for the filtered document is clustered with another tuple in the document storage structure. If the document storage structure is a tree, the tuple is clustered with another tuple in the tree if the tuple is co-located with another tuple in a 25 bucket of the tree. If the document storage structure is a hash table, the tuple for the filtered document is clustered with another tuple if the tuple for the filtered document is co-located with another tuple at an entry of the hash table. If the tuple for the filtered document is clustered with another tuple in the document storage structure, flow proceeds to block 8. If the tuple for the filtered document is not clustered with another tuple in the document storage structure, flow 30 proceeds to block 9.

In block 8, the tuple for the filtered document is clustered with another tuple in the document storage structure, and the document is determined to be similar to another document in the document collection. In one embodiment, the tuple for the document is removed from the document storage structure (if the tuple was inserted into the document storage structure), and 5 the document is not added to the document collection. If another document is to be processed, flow proceeds back to block 1.

In block 9, the tuple for the filtered document is not clustered with another tuple in the document storage structure, and the document is determined not be similar to any other document in the document collection. In one embodiment, the tuple for the document is retained 10 in the document storage structure (if the tuple was inserted into the document storage structure), and the document is added to the document collection. If another document is to be processed, flow proceeds back to block 1.

Figure 4 illustrates a flow diagram for generating the document storage structure for the document collection. Prior to processing a document according to Figure 1, the document storage structure for the document collection must be generated.

In block 31, collection statistics are obtained. As discussed above for block 2 in Figure 1, several embodiments are possible for the collection statistics. If the collection statistics are based on the document collection, a single pass of the document collection is made, and collection statistics of the document collection are ascertained. Examples of collection statistics for the document collection include the following: the number of occurrences for each token of the document collection; the number of documents in the document collection; the score (e.g., the inverted document frequency) of each token; the proximity of each token to another token in the document collection.

If the collection statistics are based on one or more other document collections, a single 25 pass of each document collection is made, and the collection statistics of the other document collections are ascertained.

In block 32, characteristics of retained tokens are determined from the collection statistics. For example, the characteristics are based on very frequent tokens and/or very infrequent tokens identified from the collection statistics. A score for each token in the 30 document collection is generated using, for example, equation (1). With the score, very infrequent tokens are identified as those tokens having a score greater than or equal to the first

threshold T1, and very frequent tokens are identified as those tokens having a score less than the second threshold T2. The thresholds T1 and T2 are used as discussed above for Figure 3. As discussed above, either one or both of the thresholds T1 and T2 can be used with the invention. As another example, one or more stop words are identified from the collection statistics and are
5 used to determine retained tokens.

In block 33, the data storage structure for the document collection is generated using the characteristics of retained tokens determined from the collection statistics in block 32. Each document in the document collection is processed according to Figure 1. For block 17, the characteristics of retained tokens from block 31 are used to retain only retained tokens in the
10 token stream. For block 8, the tuples for similar documents are clustered in the document storage structure. For block 9, additions to the document storage structure are created for each document that is not similar to another document already represented in the document storage structure. The flow diagram illustrated in Figure 1 is repeated for each document in the document collection.
15

In block 34, similar documents are removed from the document storage structure. The document storage structure is traversed, and tuples clustered together are identified. For clustered tuples, all but one of the clustered tuples are removed from the data structure, and the documents corresponding to the removed clustered tuples are removed from the document collection as being similar documents.
20

Figure 6 illustrates a plan view for a computer system for the invention. The computer system 50 includes a computer 51 for implementing the invention. The computer 51 includes a computer-readable medium 52 embodying software for implementing the invention and/or software to operate the computer 51 in accordance with the invention. As an option, the computer system 50 includes a connection to a network 53. With this option, the computer 51 is
25 able to send and receive documents from other computer systems via the network 53.

Appendix A provides exemplary pseudo-code for implementing the invention. For the pseudo-code, the SHA-1 hash algorithm is used to determine the hash value for block 4, a balanced binary tree is used as the document storage structure for block 6, and Unicode ascending order is used for block 18. Although the pseudo-code is commented, several of the
30 functions and objects of the pseudo-code are discussed next.

LoadDocDB is a function that loads the document storage structure into the computer-readable medium of the computer and creates the appropriate memory data structures. The document storage structure for the document collection contains the tuples for the document collection, namely the hash value (which is designated as the SHA-1 digest value “Digest” in the pseudo-code) and the corresponding document identifier (which is designated as “DocID” in the pseudo-code). To store the data in memory for fast access, a combination of a hash table and at least one balanced binary tree are used, as illustrated in Figure 5. Because a balanced binary tree is used to store the collisions in the hash (i.e., when two documents have the same hash value), an average access time of order $O(\log(N/B))$ is obtained, where N is the number of documents in the document collection and B is the number bins used by the hash algorithm. An object representing the document storage structure and the access method for the document storage structure are returned by the LoadDocDB function.

DigestDB is a handle to the data structure returned by the LoadDocDB function. Specifically, DigestDB is a hash table based on the SHA-1 digest and the trees of the hash values for each bin of the hash table and the corresponding access and modification methods.

GetDocsToProcess is a function that returns a list of documents to process. The return object of the GetDocsToProcess function contains a list and the access method for the data represented by the list.

SHA1 Object provides an encapsulation of the table block and the SHA-1 hash algorithm. SHA1 Object provides functions to add text blocks to the digest and to calculate the final digest value.

The Parser Object takes a document handle or stream and returns tokens until the document is traversed. The Parser is responsible for removing formatting from the document, including format tags, if any, and returning tokens. By abstracting the Parser to the basic functions, new document types can be added to the system by creating a new parser object for the new type of document.

Tree is a generic data structure. As objects are added to the Tree data structure, the objects are ordered and stored as a balanced binary tree.

Iterator is an object used as a generic pointer for traversing data structures. Iterator is used in the pseudo-code to traverse the Tree data structure. An object like the Iterator object is

common to the Standard Template Library (STL) for the C++ software language and to the collection classes of the JAVA software language.

The embodiments and examples discussed herein are non-limiting examples.

The invention has been described in detail with respect to preferred embodiments, and it
5 will now be apparent from the foregoing to those skilled in the art that changes and modifications may be made without departing from the invention in its broader aspects, and the invention, therefore, as defined in the claims is intended to cover all such changes and modifications as fall within the true spirit of the invention.

PCT/US2017/023456